Chapter 7

Assimilation Under Nonideal Conditions

In this Chapter, we present the results from a series of synthetic experiments under nonideal assimilation conditions. By nonideal conditions we generally mean that we use different inputs or parameters for the generation of the (synthetic) true fields and for the estimation algorithm. In contrast to the ideal experiments of Chapter 6, nonideal conditions allow us to subject the assimilation algorithm to harder and more realistic tests.

It is important to note that under nonideal conditions, the assimilation algorithm is not operating optimally in the strict sense of the word. Even in the ideal cases of Chapter 6, the nonlinearities in the hydrologic model and in the measurement equation lead to distortions in the normality of the error distribution functions, thus limiting our chances to conduct a fully optimal assimilation. In the nonideal scenarios, however, we explicitly specify wrong error models, which is what we do involuntarily in any field application. The resulting estimation procedure must therefore be suboptimal.

This argument appears to be contradicting our initial goal to develop a truly optimal assimilation algorithm. If the algorithm will always be suboptimal, why not use a simple and fast scheme like Optimal Interpolation right away? But recall that Optimal Interpolation is not optimal even in the ideal linear case unless we manage to come up with the exact (time-dependent) state error covariances at each update, which is quite impossible in practice. In contrast, the representer algorithm only requires the covariances of the uncertain parameters, which are much easier to specify. By using the representer algorithm, which is truly optimal in the ideal linear case, we will arguably get estimates that are much closer to the truth.

In the following experiments, we address three topics which are of major importance for the development of an operational soil moisture data assimilation system. First, we investigate the quality of the estimates when multiple assimilation windows of variable length are used (Section 7.1). This also raises the question of how we can reinitialize the variational algorithm in subsequent assimilation windows. Second, we assess the performance of the estimation when the observed precipitation data are withheld in the assimilation (Section 7.2). Finally, we investigate the influence of the soil hydraulic parameters (Section 7.3). The computational requirements of the assimilation experiments are discussed in Chapter 8.

7.1 Multiple Assimilation Windows

Maybe the biggest advantage of the variational method over sequential techniques is the implicit and thus cheap propagation of the error covariances. But in an operational setup, this is also a disadvantage. If we want to use the variational method operationally, we have to choose appropriate assimilation windows and, most importantly, we have to reinitialize these windows repeatedly. In particular, we need to calculate the posterior state error covariance at the final time of the assimilation window such that we can use it as the prior initial condition covariance for the subsequent window. Unless we can exactly compute the posterior state error covariance, we are likely to end up with a suboptimal algorithm. Unfortunately, the cost of providing accurate posterior state error covariances is prohibitive, even if we are only interested in the posterior covariances at the final time.

In practice, however, this does not necessarily mean that we are going to do very poorly even if we cannot reinitialize the assimilation windows optimally. Since the exact error covariances are implicitly propagated in the variational method, a suboptimally specified initial error covariance may evolve into the optimal error covariance after some time. This means that we should specify the assimilation window such that the observation time is near the end of the window. The wrongly specified initial error covariance will then be implicitly propagated for some time, and the error introduced at the initial time will have less impact.

This is also how we would want to proceed from a practical point of view. As soon as a brightness image becomes available, we would like to improve our estimate of the current soil moisture conditions in order to issue a forecast. We would then specify an assimilation window that starts some time in the past and ends at the current time. Of course we are not limited to having just one observation time within the assimilation window. Ideally, we would like to assimilate many past observations, although in practice we are certainly limited by the computational burden that this entails.

Since soil moisture variability is foremost governed by rain, it appears to make sense to take the precipitation history into account when choosing the assimilation windows. As strong rain events tend to wipe out the soil moisture variability at least near the surface, a natural choice is to use assimilation windows that coincide roughly with the interstorm periods. This should also make it easier to estimate the initial saturation.

7.1.1 Experiment Design

We now present the results of three experiments which differ in the choice of the assimilation windows covering the two-week experiment period. The temporal setup is illustrated in Figure 6.2. In experiment A, we use three assimilation windows that roughly coincide with the interstorm periods. In the other two experiments, we cut the two-week period into twelve short assimilation windows, each of which contains one observation time. For experiment B, we choose the assimilation windows such that the observation time is always at the end of the window. In experiment C, the assimilation windows are chosen such that the observation time is at the beginning of the window, with the exception of the first window. As discussed above, experiment B is the setup we prefer on theoretical and practical grounds.

All three experiments are based on Reference Experiment II, that is all inputs are the same unless mentioned otherwise. The reinitialization of the assimilation windows is done in

an intentionally simple and ad hoc fashion. At the beginning of each assimilation window, we use the state estimate at the final time of the previous assimilation window to derive the prior mean of the initial condition parameters as well as the shape of the initial saturation profile. For the prior covariance of the initial error we use a scaled version of the initial condition covariance of Reference Experiment II. In particular, the correlation length is always $50km$, regardless of the evolution of the system. The scaling of the covariance was found necessary to ensure convergence. Using the initial condition uncertainty of Reference Experiment II for the now shorter assimilation windows appears to lead to poor conditioning of the representer matrix, and consequently convergence could not always be achieved.

We scale the initial condition variance with a factor depending on the length of the assimilation window, or equivalently, the number of observation times within the window. Moreover, the scaling factor also depends on the number of the assimilation window. Earlier windows have a relatively higher initial condition variance. In particular, in the first experiment the three windows contain (in order) 5, 3, and 4 out of 12 observation times, and we scale the initial condition covariance of the three windows with 0.63, 0.19, and 0.17, respectively. In the other two experiments, each window contains exactly one observation time, and we scale the initial covariance of Reference Experiment II with $0.5/n_w$, where $n_w = 1 \dots 12$ indicates the number of the assimilation window. For example, the initial condition covariance of the second assimilation window is 0.25 times the initial condition covariance of Reference Experiment II.

7.1.2 Estimation of the True Fields

Figure 7.1 shows the area average root-mean-square errors (rmse) of the estimated top node saturation for the three experiments with three and twelve assimilation windows, together with the rmse of Reference Experiment II. For experiment A, the time and area average top node saturation error is 3%, and for experiments B and C the errors are 3.2% and 3.8%, respectively. These numbers compare to an error of 2.9% in Reference Experiment II. As expected, the area average error increases as the number of assimilation windows increases. With each additional assimilation window we introduce more approximations by naively reinitializing the initial saturation covariance.

Certainly the most interesting result, however, is the difference in the errors of experiments B and C. This difference illustrates the impact of the implicit error covariance propagation in the variational scheme. In experiment B, the choice of the assimilation windows lets the initial error covariance evolve dynamically for almost 24 hours before the observation time. In experiment C, by contrast, the poorly specified initial error covariance is propagated for only a couple of hours before the observation time. The update therefore relies on a crude approximation of the error covariance, which results in poorer estimates. Note that experiment C corresponds closely to an Optimal Interpolation scheme, in which the error covariances are not propagated. In summary, the estimates for Reference Experiment II and for experiments A and B are very similar. This suggests that the suboptimality introduced by the naive reinitialization is not severe, provided the assimilation windows are chosen such that the initial error covariance can evolve for at least one day.

7.1.3 Assessing the Optimality of the Estimates

We can further examine the degree of suboptimality by looking at the value of the reduced objective function. Figure 7.2 shows the reduced objective function after convergence for

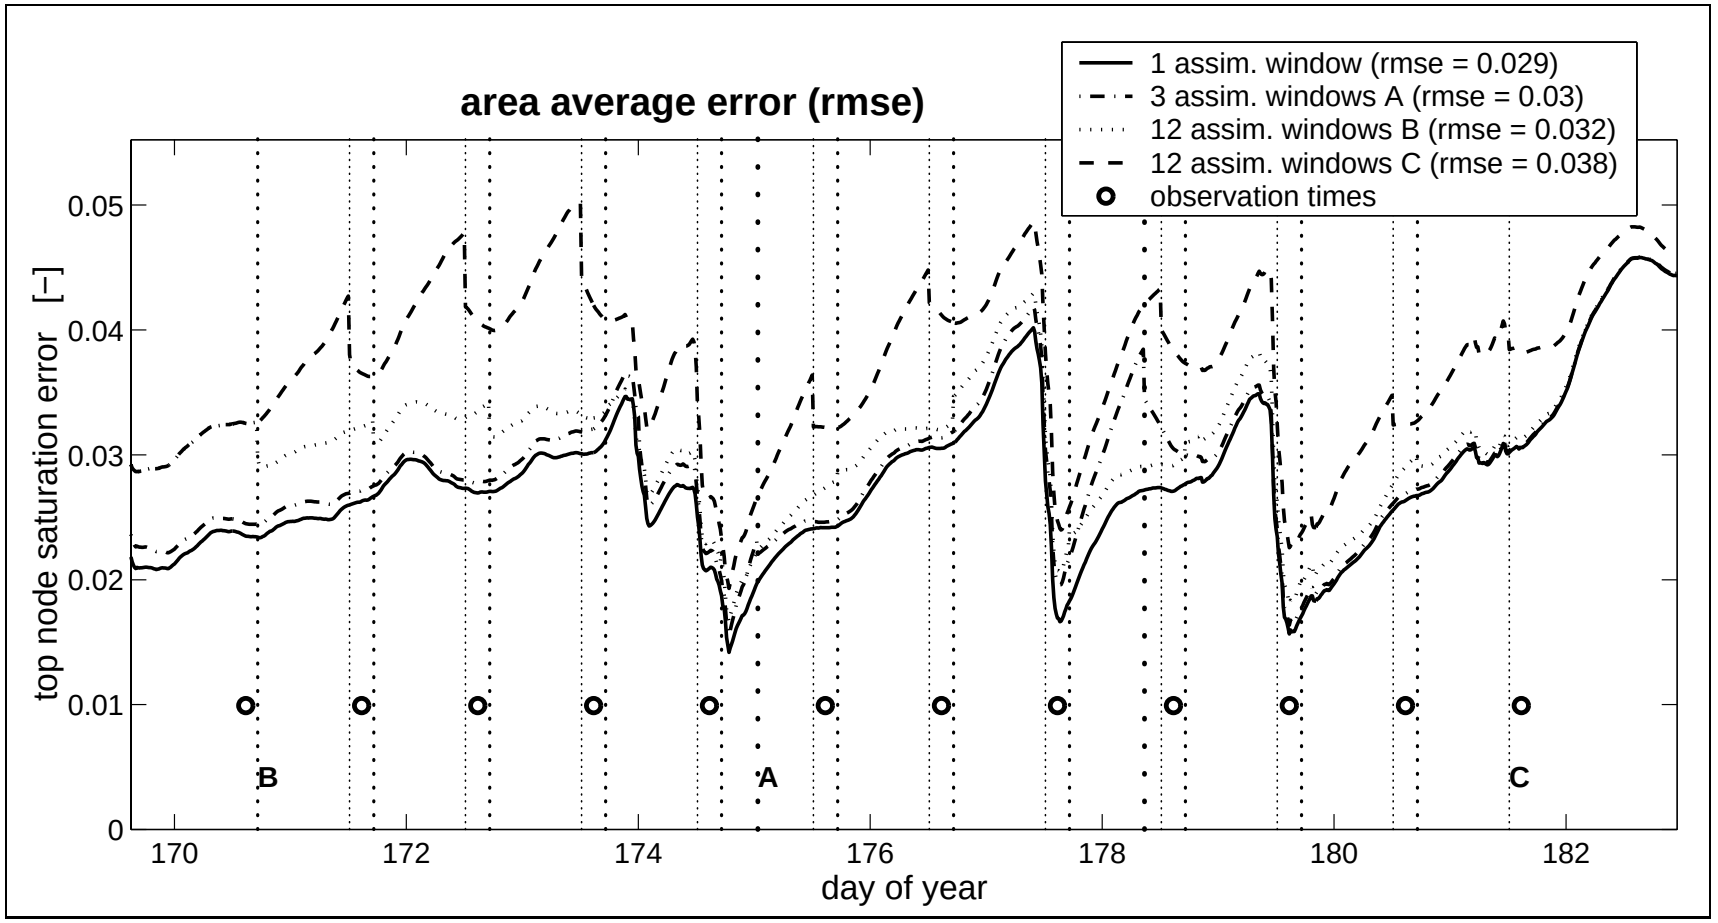

Figure 7.1: Area average errors for multiple assimilation windows. The errors in the estimated top node saturation are shown for Reference Experiment II (Figure 6.12) and the three experiments with shorter assimilation windows. The errors are in the root-mean-square sense with respect to the (synthetic) true solution. The prior root-mean-square error (rmse) is shown in Figure 6.12. In the legend we also indicate the temporal average of the area average rmse. Note that the soil moisture errors are in terms of saturation. The coarsely dotted vertical lines A delimit the three assimilation windows of the first experiment. Likewise, the more finely dotted vertical lines B and C delimit the twelve assimilation windows of the other two experiments. The observation times are indicated with circles. For experiment B the observation times are always at the end of the window, whereas for experiment C the observation times are at the beginning of the window.

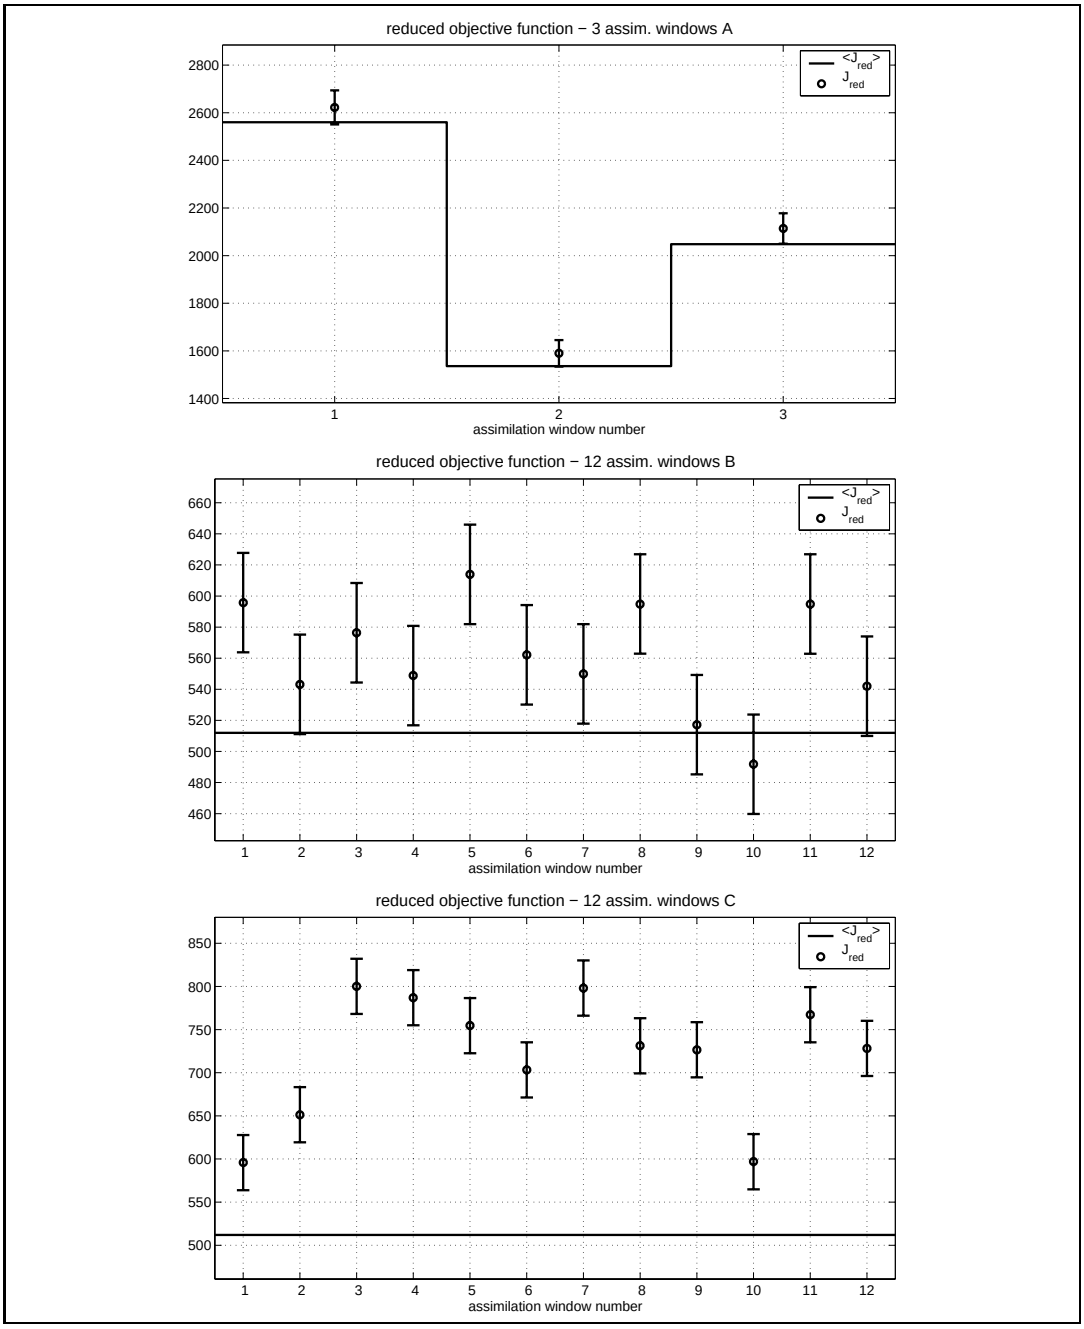

Figure 7.2: Objective function for multiple assimilation windows. The reduced objective function after convergence is plotted versus the number of the assimilation window for experiments A (top), B (middle), and C (bottom). The expected values, which equal the number of data assimilated in the window, are also shown. The error bars around the reduced objective are plus/minus one standard deviation, which is equal to the square-root of twice the number of data assimilated. For experiment A, the reduced objective function does not necessarily indicate suboptimal assimilation. Experiments B and C, however, are clearly suboptimal.

each assimilation window. We also show the expected value of the reduced objective function, which equals the number of scalar data that have been assimilated in each window. The standard deviation of the reduced objective function is indicated with error bars.

For experiment A, the reduced objective function of the three assimilation windows lies within one standard deviation from the respective expected values. Even though all three values are above the expected value, the reduced objective function does not necessarily indicate that the assimilation was suboptimal. For experiment B, eight out of twelve values of the reduced objective function are more than one standard deviation above the expected value, and all but one of the values lie above the expected value. This hints at the fact that the assimilation was not optimal. Clearly, the assimilation of experiment C must have been suboptimal.

An investigation of the posterior data residuals yields roughly the same results as for Reference Experiment II. The residuals are white in time and the hypothesis of a normal distribution cannot be rejected in almost all cases. For experiments A and B, the residuals show no obvious spatial patterns. Only the residuals of experiment C exhibit a weak spatial structure (not shown). In summary, using shorter assimilation windows and a relatively naive reinitialization does not appear to have a significant negative effect on the optimality of the algorithm, provided the assimilation windows are chosen such as to allow adequate evolution of the error covariance before observation times.

Finally, we would like to note that the computational effort for experiments A, B, and C is substantially smaller than for Reference Experiment II. For a more detailed discussion please turn to Section 8.1.3.

7.2 Assimilation without Precipitation Data

Of all model inputs, precipitation is the one parameter which dominates soil moisture conditions. At the same time, precipitation is also the input associated with the highest uncertainty. Precipitation observations from rain gauges are point measurements, and the interpolation to larger areas is notoriously ill understood. On the other hand, large-scale precipitation measurements from radar sensors are equally imprecise. The quality of the soil moisture estimates stands and falls with the accuracy of the large-scale precipitation data. It is therefore desirable to take a closer look at the sensitivity of the assimilation algorithm to the precipitation inputs.

7.2.1 Experiment Design

We now present the results of an experiment where the precipitation data are withheld from the assimilation. The experiment is based on Reference Experiment I (Section 6.1), and all inputs are the same unless otherwise mentioned. For this experiment, instead of supplying the observed precipitation time series to the assimilation algorithm, we specify zero precipitation throughout the two-week period. In order to compensate for the lack of precipitation in the assimilation, we specify certain times at which the model error in the upper moisture boundary condition has a very high variance. The times at which such model error occurs are chosen to be times at which significant area average precipitation has been observed. Such times can be regarded as precipitation indicators. This is a realistic, albeit extreme, scenario, because it is fairly easy to detect whether or not there is precipitation,

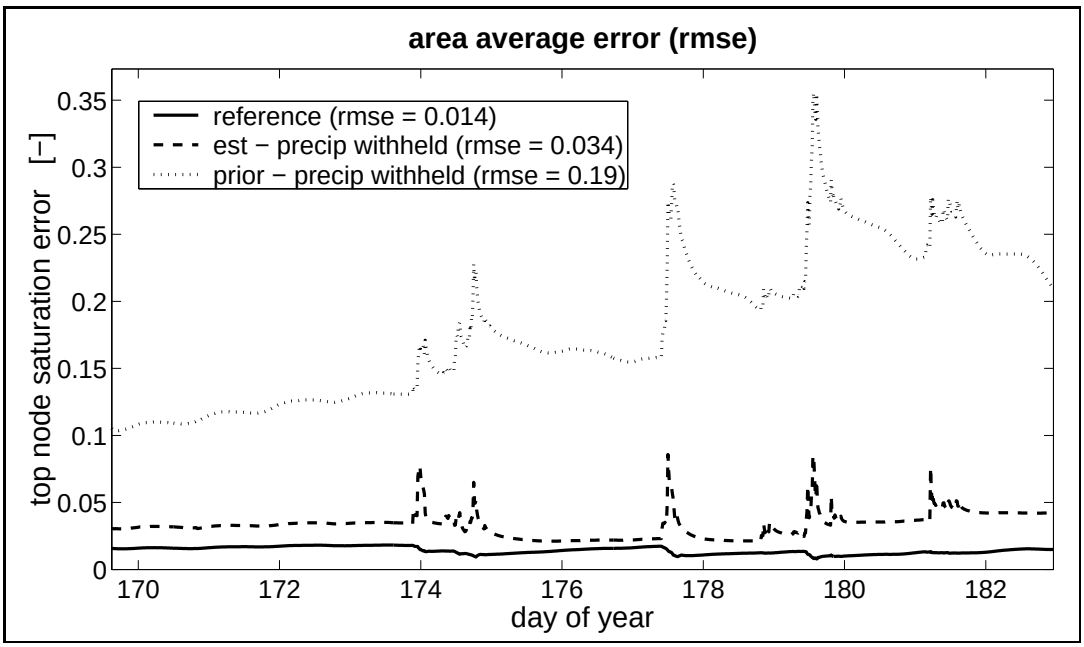

Figure 7.3: Area average errors when precipitation is withheld. The errors in the top node saturation of the estimate, the prior, and the estimate of Reference Experiment I (Figure 6.4) are shown. The errors are in the root-mean-square sense with respect to the (synthetic) true solution. In the legend we also indicate the temporal mean of the area average rmse. Note that the soil moisture errors are in terms of saturation (see Figure 6.4). Even when quantitative precipitation data are withheld entirely, the algorithm can estimate soil moisture to good accuracy.

but it is rather difficult to observe the rain event quantitatively. In practice, one would of course use whatever quantitative precipitation information is available.

Note that the intermittent model error is nonstationary, and in practice we then have to use a white noise error model to make the noise update computationally feasible (Section 4.8). When precipitation is indicated, we set the spatially uniform standard deviation of the model error in the moisture flux boundary condition equal to $86.4 mm/d$, which is a typical area average rain rate at $15 min$ resolution. This is the case for 70 out of 1280 time steps. At all other times, the white component of the model error has zero variance.

Note that we did not tune the error model to improve the performance of the estimation algorithm. Moreover, the above error model is of course not the only choice. Many other error models can be thought up to compensate for the lack of precipitation data in the assimilation. For example, the model error variance could be horizontally distributed, which may better capture the strong gradient in rainfall across the domain. Also note that with the above error model, "negative precipitation" could result from the estimation, that is the model error could lead to unrealistically high evapotranspiration rates. To prevent this from happening, the model error could be formulated from the start with a logarithmic transform.

7.2.2 Estimation of the True Fields

Figure 7.3 shows the area average top node saturation errors with respect to the (synthetic) true fields when quantitative precipitation data are withheld from the assimilation. For comparison, we also show the best estimate of the ideal assimilation run of Section 6.1 (Figure 6.4). Even when all quantitative precipitation information is withheld from the assimilation, we can estimate the top node saturation to within 3.4% in saturation terms, compared to 1.4% under ideal conditions. Unlike in the case of the ideal Reference Experiment I (Figure 6.4), the prior error in this experiment increases over the two-week assimilation window. Since the prior solution does not contain any rain events, the prior saturation is governed by a single, two-week long drydown. For longer assimilation windows, the prior saturation error reaches a plateau.

Around precipitation events, the average error in the top node saturation increases somewhat. This follows naturally from the fact that we cannot resolve the temporal structure of the events from brightness data that are available only once daily. To illustrate this point, we plot in Figure 7.4 the observed precipitation and the corresponding model error estimates for three different pixels during the three major precipitation events of the two-week assimilation window. The three pixels shown in Figure 7.4 are in the southwestern corner (pixel 100), the center (pixel 398), and the northeastern corner (pixel 412) of the domain. For the three pixels, the cumulative observed precipitation over the entire two-week period is $1.6cm$, $1.5cm$, and $4cm$, respectively. The cumulative model error estimates are $1.6cm$, $0.9cm$, and $2.8cm$. The area average cumulative precipitation is $2.8cm$, and the corresponding area average cumulative model error estimate is $1.6cm$.

By using the precipitation indicators we naturally get the overall timing of the storms right. In addition, we also get reasonable estimates of the volume of the storms from the brightness observations. It is clear, however, that the detailed temporal structure of the storms eludes us. Moreover, the model error estimates are generally lower than the observed precipitation. To understand this, recall that we assimilate brightness data only once a day. In this case, the nonlinearities in the infiltration and exfiltration processes defy a more accurate estimate of the volume of the storms. Since the true precipitation is heavier than the model error estimates, the soil actually gets wetter than estimated, but this also leads to stronger evaporation and possibly runoff. If the observations are available only some time after the storm, the difference between the true and the estimated soil saturation at the observation time is then much smaller than the difference between the model error estimate and the observed precipitation at the time of the storm, and we cannot distinguish between the two scenarios.

Note that the increase in the model error estimate with time during blocks of nonzero estimates is a direct consequence of the whiteness of the model error. If the temporal correlation of the model error is white, the model error estimates are in essence scaled version of the adjoint variables. The adjoint variables, in turn, decay backwards in time, as can be seen for example in Figure 6.10.

7.2.3 Reduced Objective Function and Computational Effort

Figure 7.5 shows the reduced objective function during the iteration. After convergence, the reduced objective function is 10,907, which is not compatible with an expected value of 6,144. Obviously, our choice of error model does not capture the real error introduced by withholding all quantitative precipitation information. In fact, the above experiment is only

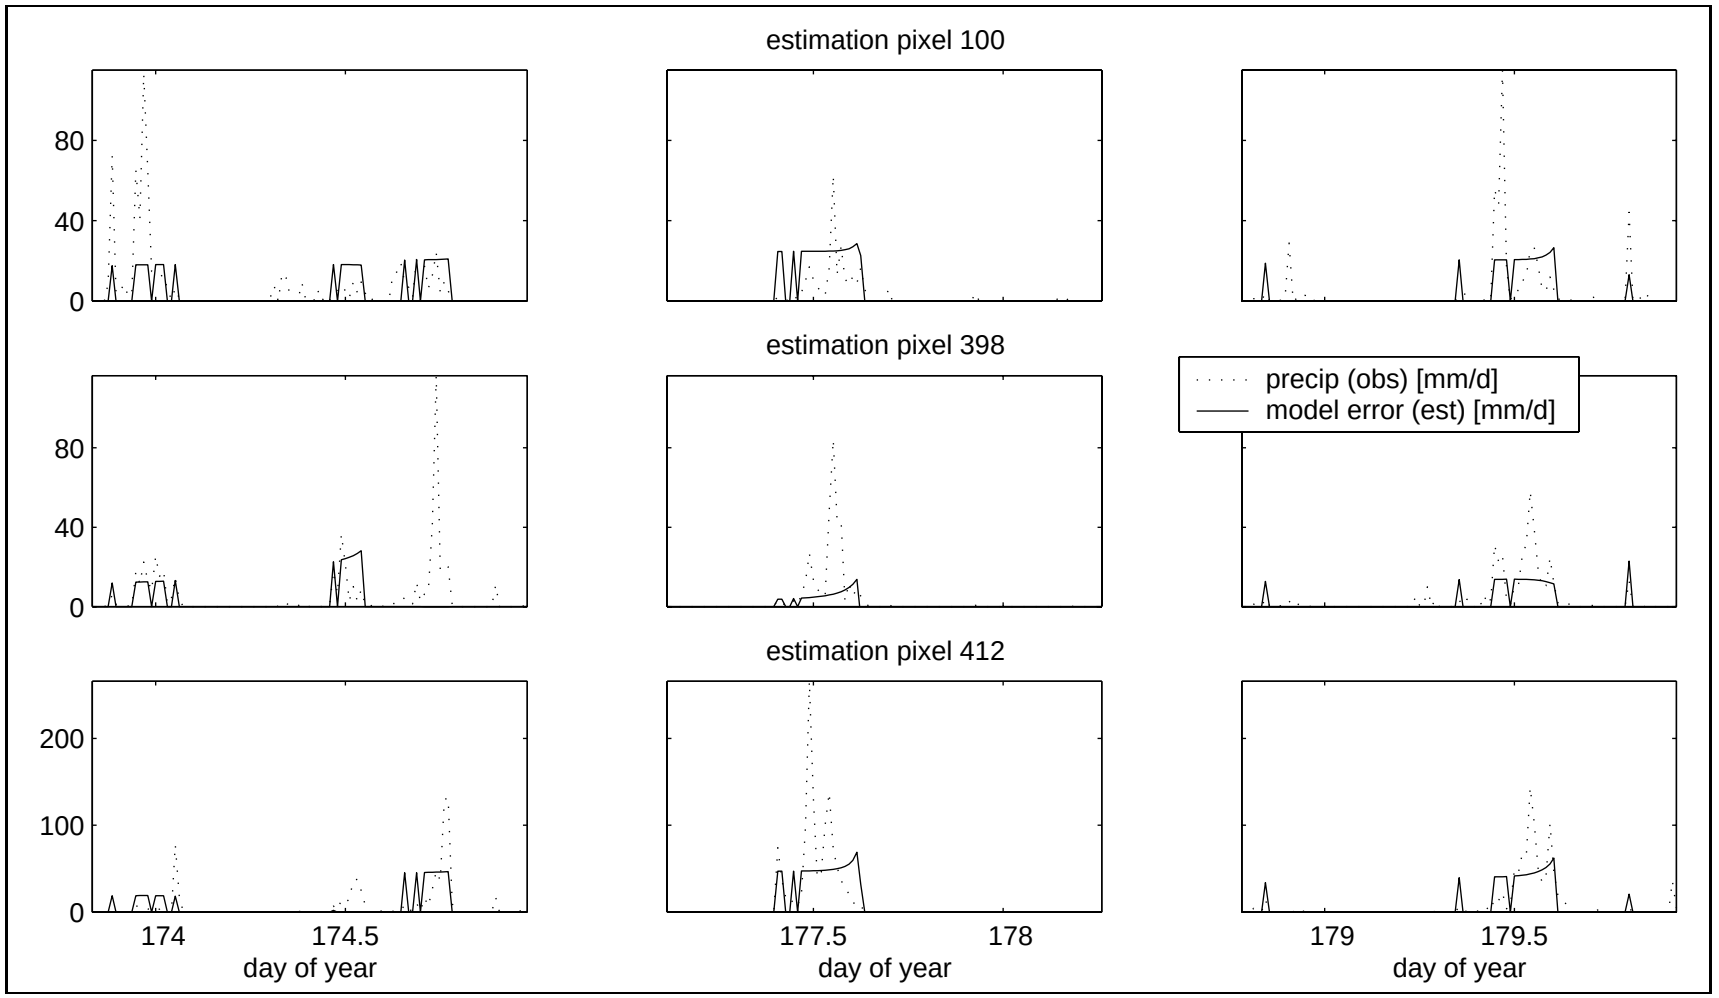

Figure 7.4: Model error estimates when precipitation is withheld. Model error estimates for three different pixels during the three major precipitation events. The pixels 100, 398, and 412 are located in the southwestern corner, the center, and the northeastern corner of the domain, respectively. By using precipitation indicators, we supply the overall timing of the storms to the estimation algorithm. By assimilating brightness observations, we get reasonable estimates of the volume of the storms. To better compare the model error estimates to the observed precipitation, which is defined as a positive quantity, we plot the negative of the model error estimate.

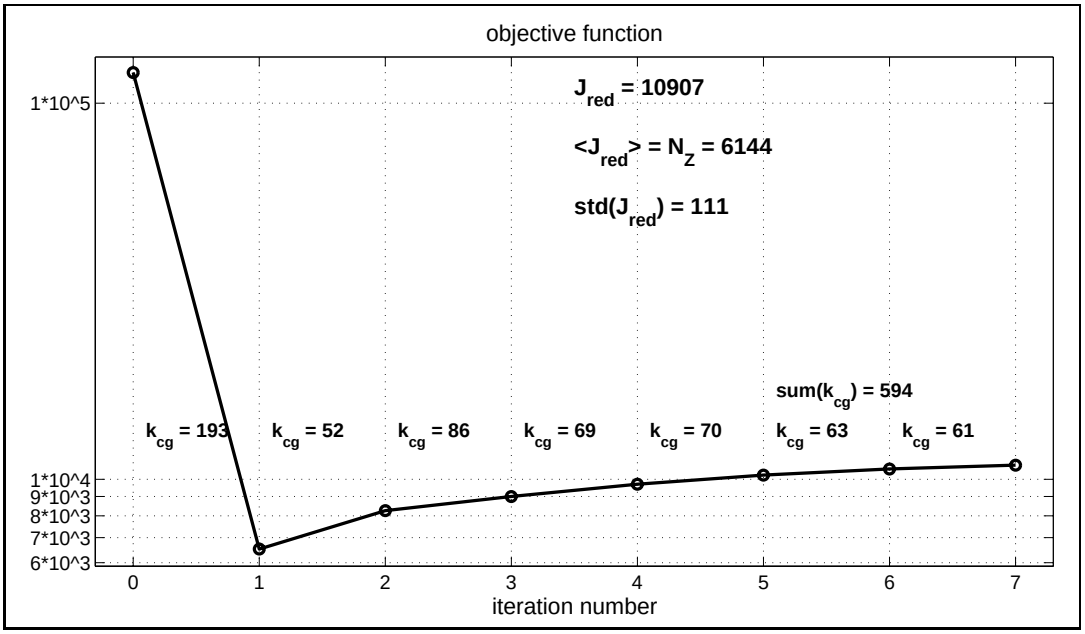

Figure 7.5: Objective function versus iteration number when precipitation is withheld. The reduced objective function after convergence is 10,907. The number of data points is 6144, which is also the expected value of the reduced objective function. Obviously, the estimates do not pass the hypothesis test on the value of the reduced objective function. The values of k_{cg} indicate the number of linear combinations of representer functions that needed to be evaluated during the conjugate gradient iteration of the indirect representer approach (Chapter 8).

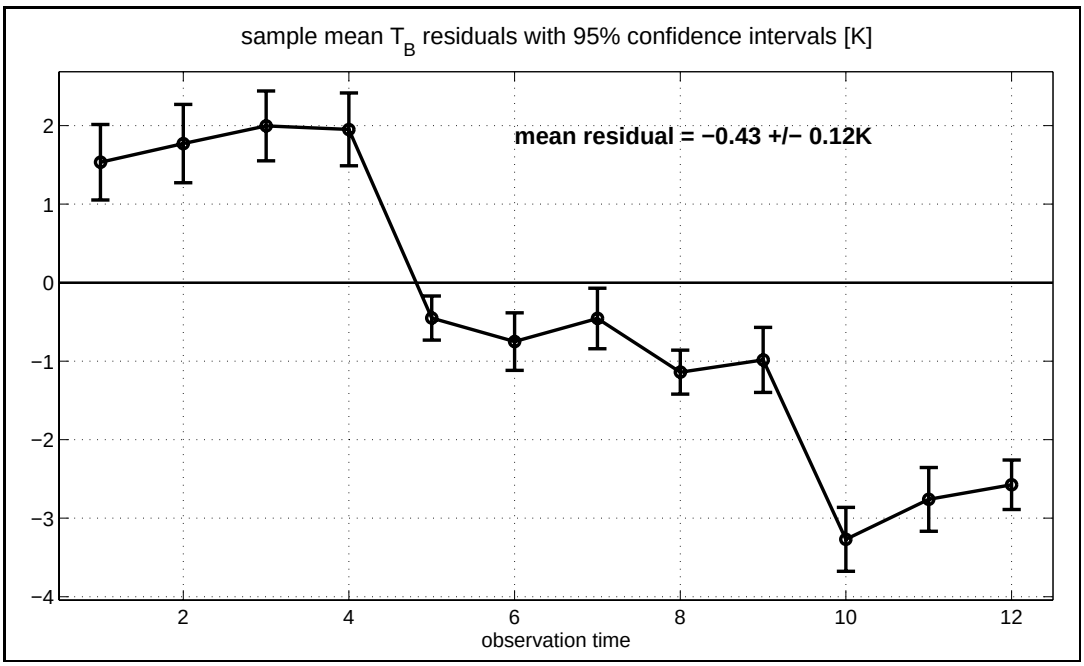

Figure 7.6: Sample mean of the posterior data residuals when precipitation is withheld. The mean residuals for the twelve images that have been assimilated are shown together with the 95% confidence intervals. Very obviously, the mean value of the posterior data residuals differs from zero.

a crude approximation of a much more sophisticated optimal assimilation procedure. Ideally, if the precipitation inputs are considered very uncertain, the true precipitation should be estimated by assimilating the rain data into a suitable model of precipitation processes. In the future, the current soil moisture assimilation algorithm may be augmented to also estimate precipitation.

Interestingly, the computational effort for this experiment is only about two thirds of the computational burden of Reference Experiment I. This stems from the fact that we changed the prior statistics in order to compensate for the lack of precipitation data. For a more detailed discussion see Section 8.1.2.

7.2.4 Posterior Data Residuals

A closer look at the data residuals sheds more light on the suboptimal nature of the assimilation in this experiment. We first examine the mean of the residuals (Figure 7.6). Unsurprisingly, the mean of all residuals with a 95% confidence interval is $-0.43 \pm 0.12K$, which does not include zero. Similarly, none of the individual residual brightness images has a mean whose 95% confidence interval includes zero. Figure 7.7 shows the standardized posterior data residuals for each of the twelve brightness images that have been assimilated. Obviously, some of the residual images show spatial structure. This is especially true for observation time 10, which happens to be during a major rain event in the northern half of the domain.

Figure 7.8 shows the sample cumulative distribution function (cdf) for two representative

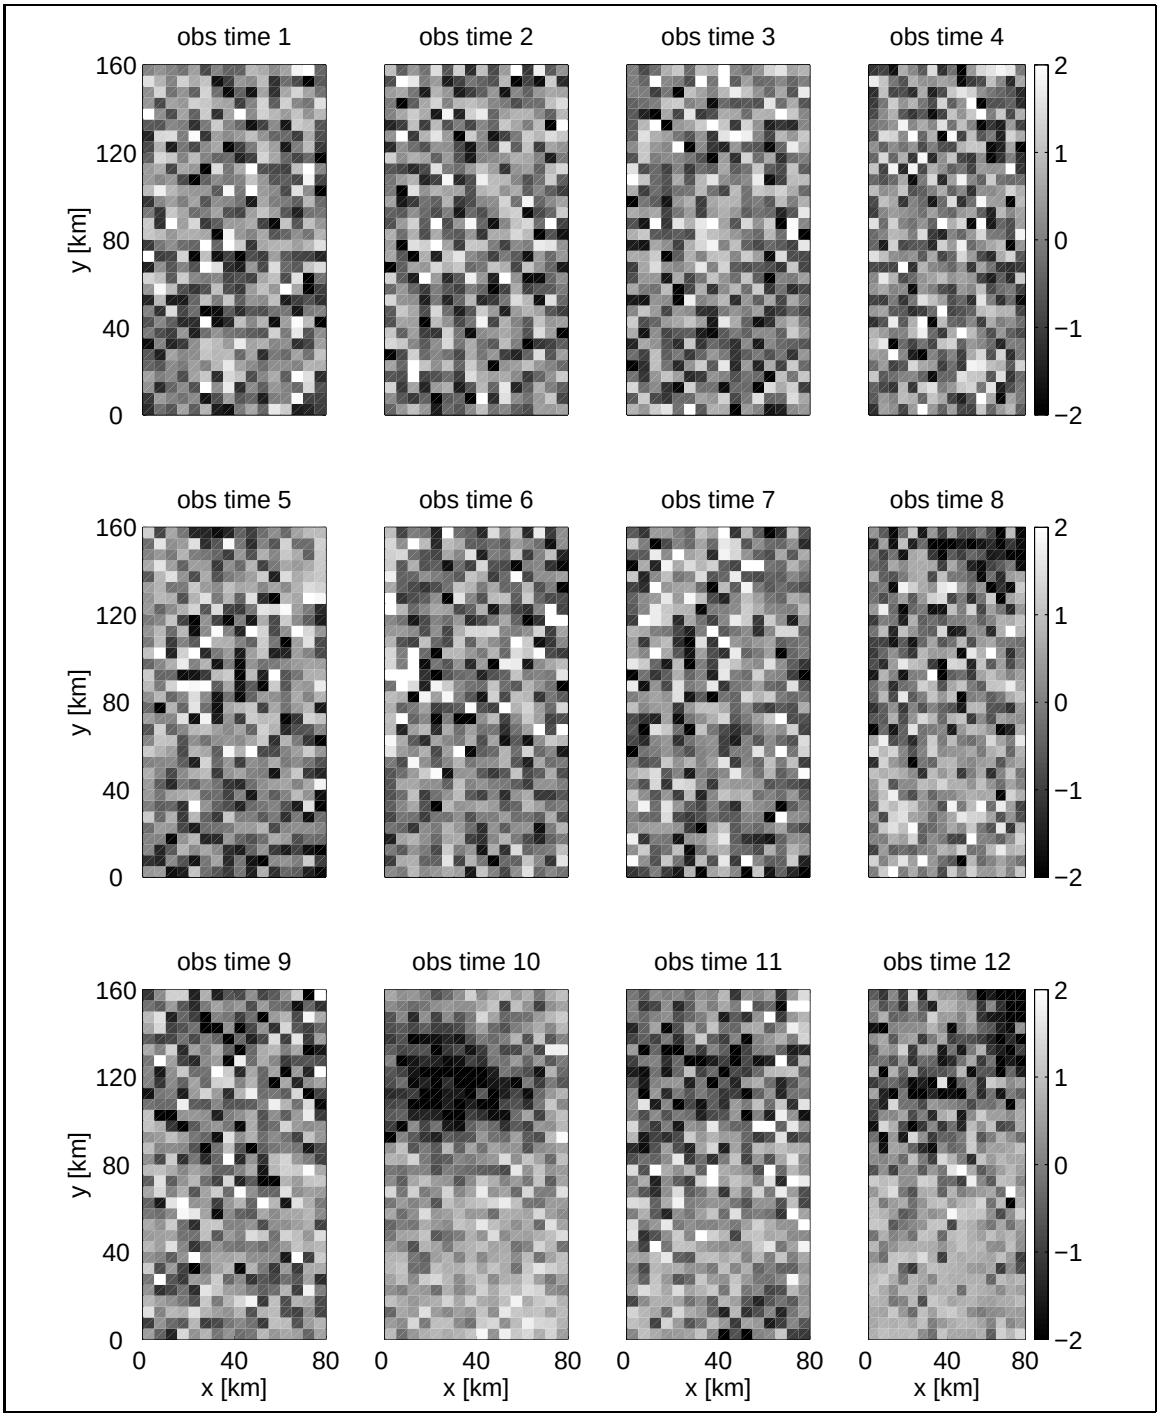

Figure 7.7: Standardized posterior data residuals when precipitation is withheld for the twelve brightness images that have been assimilated. The residuals of each image are standardized with the sample mean and standard deviation of the corresponding observation time. Some of the residual images show an obvious spatial structure, which indicates that the estimation algorithm does not work optimally.

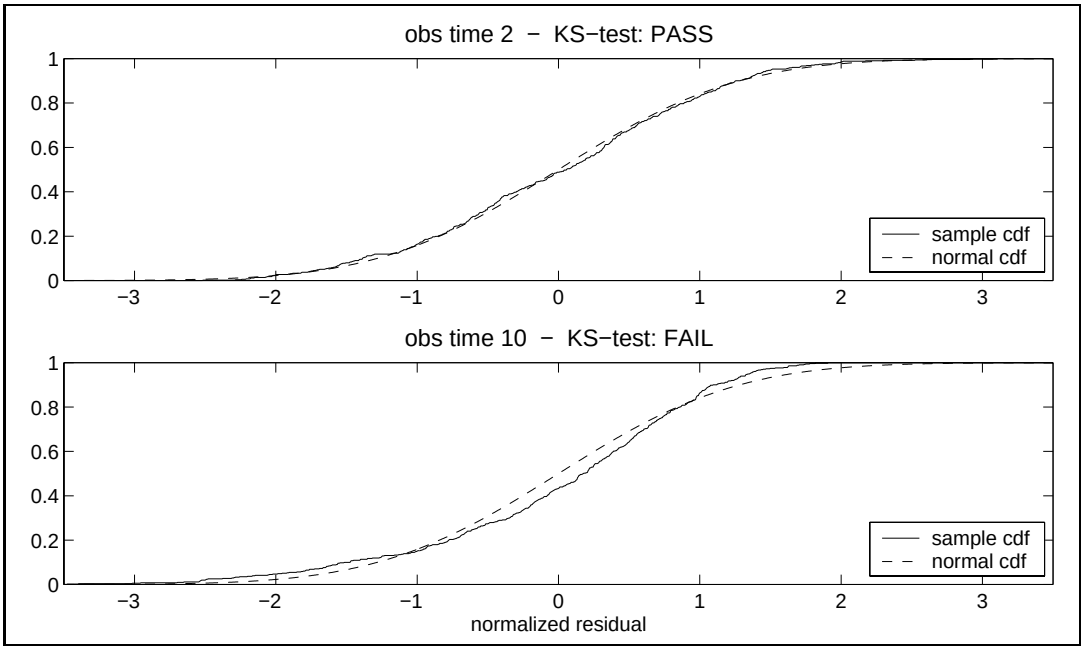

Figure 7.8: Sample cumulative distribution function (cdf) of the standardized posterior data residuals when precipitation is withheld for two of the twelve brightness images that have been assimilated (solid line). Also shown is the theoretical cumulative distribution function of the standard normal distribution (dashed line). For observation time 2 (upper panel), the residuals are close to normal and pass the Kolmogorov-Smirnov test for normality at a significance level of 5%. For observation time 10, however, the sample cdf is far from normal. The same is true for four other observation times.

residual images. For seven of the twelve observation times, the sample cdf is close to normal and passes the Kolmogorov-Smirnov test at a significance level of 5%. A representative example is observation time 2, which is shown in the upper panel of Figure 7.8. The lower panel shows an example of the five observation times for which the residuals are far from normally distributed. The examination of the brightness residuals thus corroborates that the assimilation was not optimal.

7.3 Assimilation with Poor Soil Hydraulic Parameters

Inevitably, the success of the estimation procedure depends on the quality of the model parameters that need to be specified. Many of these inputs, in particular the soil hydraulic parameters, are rather poorly known. Ideally, one would of course estimate the uncertain model parameters. Even though the data assimilation algorithm formulated in Chapter 2 is very general and provides for the estimation of model parameters, we have not implemented this feature in the synthetic experiments of this thesis for two reasons. First, the state estimation problem as implemented is already very complicated, and it is certainly wise to be conservative when specifying the uncertain inputs in a first application. Second, any hydrologic model that is to be used in an operational assimilation package had better be well calibrated. Estimating already calibrated parameters may make sense when one tries to improve the stability of the assimilation algorithm, but it is unlikely to be the most pressing problem in an operational context. Moreover, the parameter estimate adds significantly to the computational burden.

7.3.1 Experiment Design

Nevertheless, it is important to understand the sensitivity of the assimilation algorithm to poorly or wrongly specified model parameters. To address this issue, we have conducted a synthetic experiment in which we use different soil hydraulic parameters for the generation of the (synthetic) true fields and for the estimation. The setup of the experiment, including the (synthetic) true solution, is identical to Reference Experiment I of Section 6.1, with one exception. For the estimation, we change the soil hydraulic parameters of the land surface model by assigning the values from soil texture classes that have been randomly sampled from the existing soil texture classes. The resulting soil texture map is shown in Figure 7.9. Compared to the original map (Figure 5.3), 361 out of the 512 pixels differ in their texture classes.

We choose to randomly sample from existing soil texture classes for two reasons. First, the procedure guarantees that we only work with calibrated and tested input parameters. Unfortunately, the stability of the Richards' equation solver is fairly sensitive to the soil hydraulic parameters. Moreover, the scenario is realistic because in general we will have a good idea what soil texture classes occur in any given area, even though we may not accurately know their spatial distribution.

7.3.2 Estimation of the True Fields

Figure 7.10 shows the area average errors in the top node saturation of the prior and the estimate for the texture sensitivity experiment. Also shown is the error of the estimate from the ideal setup of Reference Experiment I (Figure 6.4). Even for wrongly specified soil

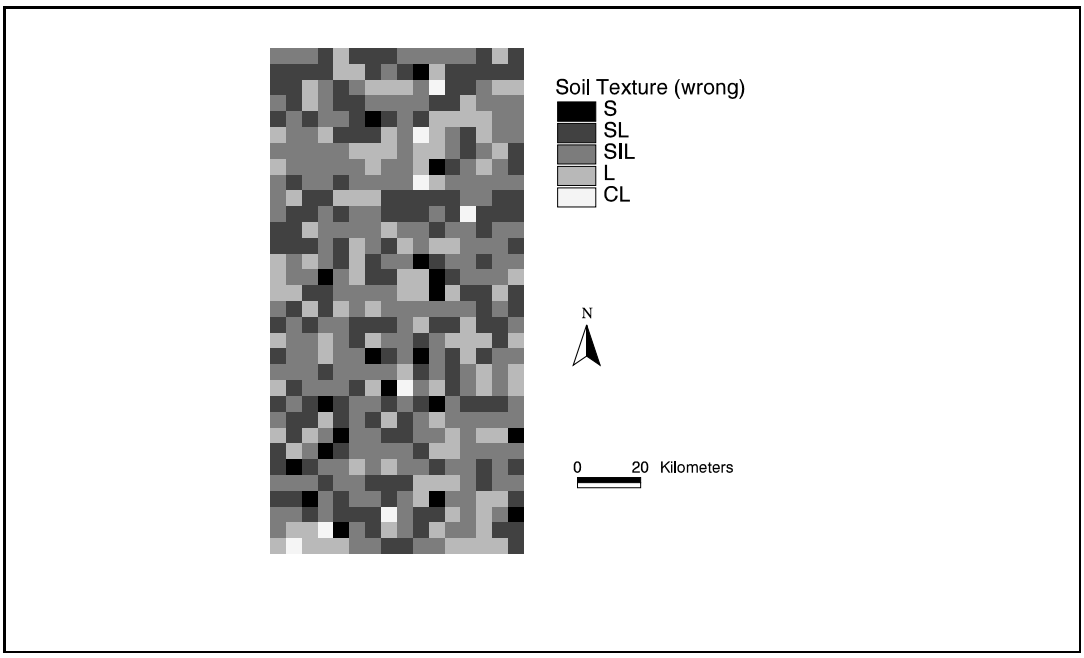

Figure 7.9: Soil texture classes for the texture sensitivity experiment. This "wrong" soil texture map has been derived by randomly subsampling from the original soil texture map (Figure 5.3). Compared to the original map, 361 out of the 512 pixels differ in their texture classes.

hydraulic parameters, the assimilation algorithm manages to estimate soil moisture satisfactorily. During precipitation events, however, the error increases. This is directly related to the fact that we supply wrong soil hydraulic parameters to the assimilation algorithm. Most of the infiltration happens during rain events, and exfiltration via evapotranspiration is highest just after the events. But with wrong soil hydraulic parameters, we cannot model the (synthetic) true moisture fronts accurately. This leads to higher errors during and after rain events.

Note that using the wrong soil hydraulic parameters leads to a prior error (Figure 7.10) which is only slightly higher than the error of Reference Experiment I (Figure 6.4), in which the true texture classes have been used. In Reference Experiment I, we already assume a very big uncertainty in the initial condition, and the prior rmse is about as big as it can be under the conditions of the experiment. Since we already know next to nothing about the initial condition distribution, the prior error does not increase much if we use the wrong soil hydraulic parameters.

7.3.3 Assessing the Optimality of the Estimates

Finally, Figure 7.11 shows the reduced objective function for the texture sensitivity experiment. The converged value of the reduced objective function is 7023, which is almost eight standard deviations above the expected value. This follows naturally from the fact that we use the same error statistics for the generation of the (synthetic) true solution and for the estimation, even though we changed the soil hydraulic parameters in the assimilation.

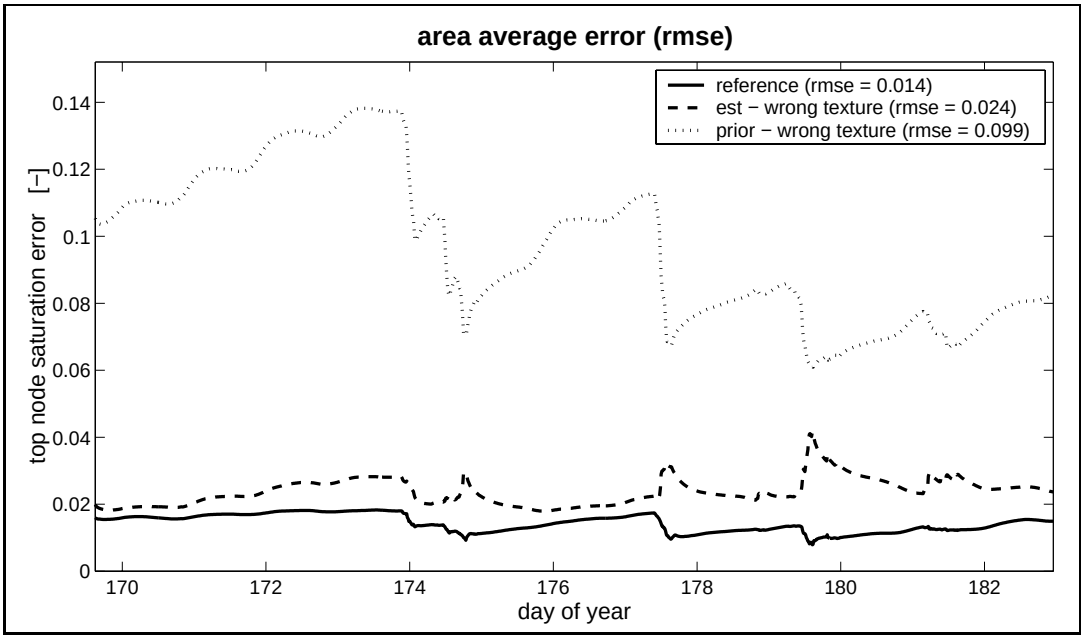

Figure 7.10: Area average errors for the texture sensitivity experiment. The errors in the top node saturation of the estimate, the prior, and the estimate of Reference Experiment I (Figure 6.4) are shown. The errors are in the root-mean-square sense with respect to the (synthetic) true solution. In the legend we also indicate the temporal mean of the area average rmse. Note that the soil moisture errors are in terms of saturation.

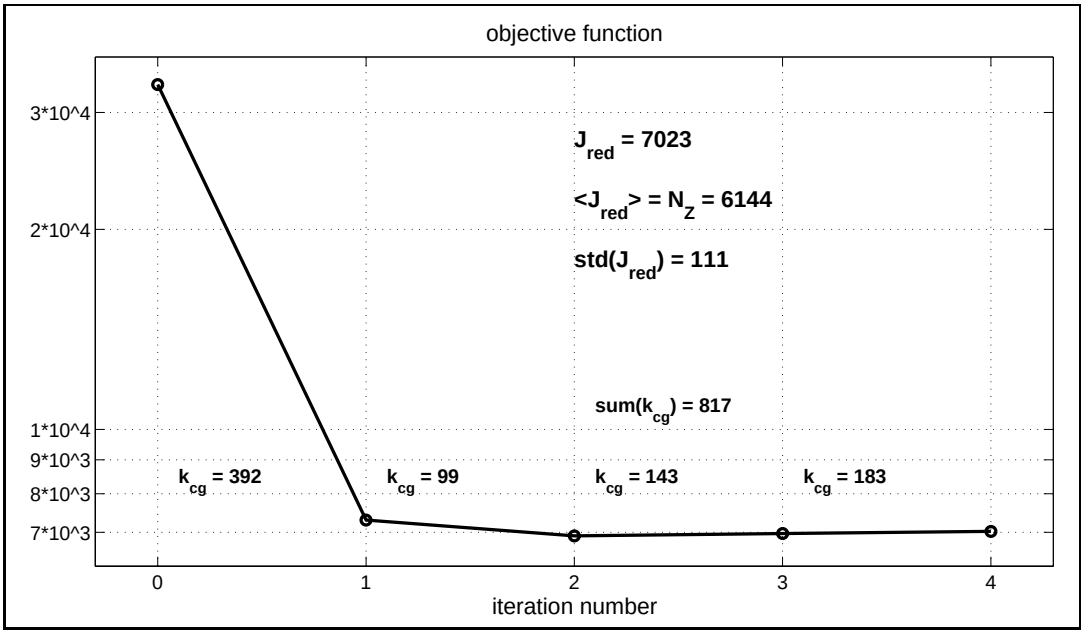

Figure 7.11: Objective function versus iteration number for the texture sensitivity experiment. The reduced objective function after convergence is 7023. The number of data points is 6144, which is also the expected value of the reduced objective function. Obviously, the estimates do not pass the hypothesis test on the value of the reduced objective function. The values of k_{cg} indicate the number of linear combinations of representer functions that needed to be evaluated during the conjugate gradient iteration of the indirect representer approach (Chapter 8).

The change in model parameters leads to additional errors in the model, which are not adequately described with the original error statistics. This confirms that the assimilation is not operating optimally in this nonideal case.

On the other hand, a close look at the posterior data residuals yields qualitatively the same results as for Reference Experiment I (Figures 6.7, 6.8, and 6.9). The raw mean for all residuals with a 95% confidence interval is $0.13 \pm 0.13K$. For all but one of the residual brightness images we find a mean whose 95% confidence interval includes zero. There is no indication that the residuals are correlated in either space or time. Moreover, the sample cumulative distribution function for all but one of the residual images passes the Kolmogorov-Smirnov test for normality. In summary, this indicates that the estimation process is at least close to optimal.